Experiment 6a - Bidirectionial Forwarding Detection

DE-CIX Academy

Version 1.0

1 Network Setup

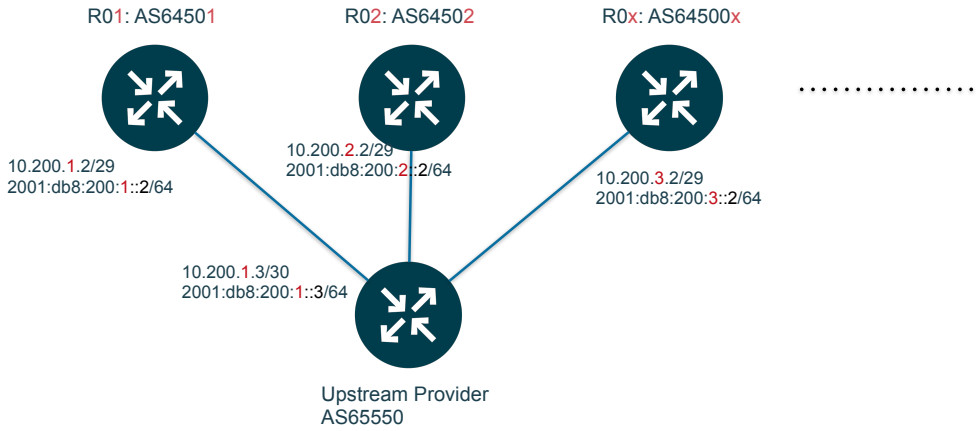


Figure 1: Network Setup

Figure 1 shows the network topology for this experiment. All students' routers connect to one upstream AS65550 (also running FRRouting in a Docker container)

2 Introduction

Timers in BGP are pretty conservative - it can take up to three minutes (with default timer values) for BGP to detect that a remote BGP speaker is no longer available *if* the network interface stays up.

BFD solves this problem by fast-sending keepalive packets to your neighbor. As this can be implemented in hardware on a linecard, this is way more efficent than setting the BGP timers to shorter intervals.

3 Task: Implement BFD to your upstream

- Configure BFD with your upstream as neighbor
- Add a statement to BGP to take down your neigbor if BFD fails.
-

4 Config statements

4.1 BFD

BFD enters the *BFD* configuration context. All further comands are within that context:

peer *a.b.c.d* creates a BFD peer to listen and talk BFD to. Also enters configuration context of that peer.

4.2 BFD and BGP

neigbor *a.b.c.d* bfd tells BGP to immediately take down the BGP session if BFD towards this neigbor fails and restart the session if BFD goes back up.

5 Show commands

5.1 BFD

show bfd peers brief gives a list of all BFD peers

show bfd peers counters gives a slightly more verbose list including all counter values

sho bdf peers very verbose list of all BFD peers

sho bdf peer a.b.c.d shows one BDF peer

A Upstream configuration

Configuration of the “upstream provider”:

```
hostname upstream
!
router bgp 65550
  no bgp default ipv4-unicast
  bgp disable-ebgp-connected-route-check
  neighbor customer peer-group
  neighbor customer bfd
  neighbor customer-v6 peer-group
  neighbor customer-v6 bfd
  neighbor 10.200.x.2 remote-as 645xx
  neighbor 10.200.x.2 peer-group customer
  neighbor 2001:db8:200:x::2 remote-as 645xx
  neighbor 2001:db8:200:x::2 peer-group customer-v6
  address-family ipv4 unicast
    network 198.51.100.0/24
    redistribute static route-map redistribute-static
    neighbor customer activate
    neighbor customer soft-reconfiguration inbound
    neighbor customer route-map customer-in in
    neighbor customer route-map customer-out out
  exit-address-family
!
  address-family ipv6 unicast
    neighbor customer-v6 activate
    neighbor customer-v6 soft-reconfiguration inbound
    neighbor customer-v6 route-map customer-in in
    neighbor customer-v6 route-map customer-out out
  exit-address-family
exit
bfd
  peer 10.200.x.2
```

```
    exit
    peer 2001:db8:200:x::2
    exit
  exit
end
```

B Final Router Configuration

Router config with BFD solution applied

```
  hostname rXX
  !
  interface loopback0
    ip address 172.16.x.1/32
   !
  ip prefix-list my-networks permit 172.16.x.0/24
  ip prefix-list my-networks-all permit 172.16.x.0/24 le 32
  !
  interface eth1
  !
  interface eth2
   ip address 10.200.x.2/29
   ipv6 address 2001:db8:200:x::2/64
  !
  route-map upstream-in permit 100
   set local-preference 10
  !
   route-map upstream-out permit 50
   match ip address prefix-list my-networks
  !
  route-map upstream-out deny 100
  !
  !
  route-map peering-in permit 100
   set local-preference 1000
  !
  route-map peering-out deny 100
  !
  router bgp 645xx
   no bgp default ipv4-unicast
   bgp ebgp-requires-policy
   neighbor upstream peer-group
   neighbor upstream bfd
   neighbor upstream-v6 peer-group
   neighbor upstream-v6 bfd
   address-family ipv4 unicast
    network 172.16.x.0/24
    redistribute static route-map redistribute-static
    neighbor upstream activate
    neighbor upstream route-map upstream-in in
    neighbor upstream route-map upstream-out out
    neighbor upstream soft-reconfiguration inbound
   exit-address-family
   address-family ipv6 unicast
    network 2001:db8:x::0/48
    redistribute static route-map redistribute-static
    neighbor upstream-v6 activate
    neighbor upstream-v6 route-map upstream-in in
```

```
   neighbor upstream-v6 route-map upstream-out out
   neighbor upstream-v6 soft-reconfiguration inbound
  exit-address-family
  neighbor 10.200.x.3 remote-as 65550
  neighbor 10.200.x.3 peer-group upstream
  neighbor 2001:db8:200:x::3 remote-as 65550
  neighbor 2001:db8:200:x::3 peer-group upstream-v6
 !
 exit
 bfd
   peer 10.200.1.3
   exit
   peer 2001:db8:200:1::3
   exit
   !
  exit
 end
```

C Slides

How fast is BGP without BFD?

- upstream will shut the interface to your router
- check how long it takes for the BGP session to go down
- just do a **show bgp summary**

Some other show commands:

- **show bgp ipv4 neighbor 10.200.x.3** - check the timer values
- **show bgp ipv6 neighbor 2001:db8:200:x::3**

Configure BFD

- BFD is an IP protocol
- Simplest configuration: Give target IP address

```
conf t

bfd
peer 10.200.x.3
peer 2001:db8:200:x::3
```

Some show commands:

- **show bfd peers brief** - list all BFD peers
- **show bfd peers** - list all BFD peers (verbose)
- **show bdf peer <address>** - show one BDF peer

Configure BGP for BFD

- We need to tell BGP to use BFD
- This is even more simple

```
conf t

router bgp 645xx
  neighbor upstream bfd
  neighbor upstream-v6 bfd
```

Some show commands:

- **show bgp neighbor 10.200.x.3** - check BFD status at the end

Now lets shut down the interface at the upstream and have check your side:

- **show bgp summary** - check how long it takes